

HEWLETT-PACKARD COMPANY

Intellectual Property Administration
P. O. Box 272400
Fort Collins, Colorado 80527-2400

PATENT APPLICATION **A**ATTORNEY DOCKET NO. 10004565

IN THE U.S. PATENT AND TRADEMARK OFFICE
Patent Application Transmittal Letter

ASSISTANT COMMISSIONER FOR PATENTS
Washington, D.C. 20231

Sir:

Transmitted herewith for filing under 37 CFR 1.53(b) is a(n): ☒ Utility ☐ Design
☒ original patent application,
☐ continuation-in-part application

INVENTOR(S): **John Wilkes and Theodore M. Wong**TITLE: **EXCLUSIVE CACHING IN COMPUTER SYSTEMS**

Enclosed are:

☒ The Declaration and Power of Attorney. ☒ signed ☐ unsigned or partially signed
☒ 3 sheets of drawings (one set) ☐ Associate Power of Attorney
☐ Form PTO-1449 ☐ Information Disclosure Statement and Form PTO-1449
☐ Priority document(s) ☐ (Other) _____ (fee \$ _____)

CLAIMS AS FILED BY OTHER THAN A SMALL ENTITY				
(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) TOTALS
TOTAL CLAIMS	19 — 20	0	X \$18	\$ 0
INDEPENDENT CLAIMS	2 — 3	0	X \$78	\$ 0
ANY MULTIPLE DEPENDENT CLAIMS	0		\$260	\$ 0
BASIC FEE: Design (\$310.00); Utility (\$690.00)				\$ 690
TOTAL FILING FEE				\$ 690
OTHER FEES				\$
TOTAL CHARGES TO DEPOSIT ACCOUNT				\$ 690

Charge \$ 690 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16, 1.17, 1.19, 1.20 and 1.21. A duplicate copy of this sheet is enclosed.

"Express Mail" label no. EJ324372530USDate of Deposit 8-17-2000

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to: Assistant Commissioner for Patents, Washington, D.C. 20231

By Paul H. Horstmann
Typed Name: **Paul H. Horstmann**

Respectfully submitted,

John Wilkes and Theodore M. WongBy Paul H. Horstmann**Paul H. Horstmann**

Attorney/Agent for Applicant(s)

Reg. No. **36,167**Date: **8-17-2000**Telephone No.: **(415) 602-1721**

BACKGROUND OF THE INVENTION

Field of Invention

5 The present invention pertains to the field of computer systems. More particularly, this invention relates to exclusive caching in a computer system.

Art Background

10 A typical computer system includes one or more host systems and one or more storage systems. A storage system usually provides relatively large-scale non-volatile storage of information which may be accessed by a host system. A host system typically accesses a storage system by performing
15 write and read operations to and from the storage system via a communication path between the host and storage systems.

20 A typical host system includes a host processor and a host cache. A typical host cache temporarily holds information obtained from a storage system and provides the host processor with relatively fast access to information held in the host cache. A storage system commonly includes a storage medium and
25 a storage system cache. A typical storage system cache temporarily holds information obtained from the storage medium and provides a host system with relatively fast access to the information contained in the storage system cache.

30 In prior host systems, the host cache is usually managed in a manner designed to increase the likelihood that the information accessed by the host

processor will be contained in the host cache.
Similarly, the storage system cache in prior storage
systems is usually managed to increase the likelihood
that the information accessed by a host system will
5 be contained in the storage system cache. As a
consequence, the same information, i.e. the
information likely to be accessed by a host
processor, often accumulates in both the host cache
and the storage system caches. Unfortunately, such
10 duplication of information reduces the effective
storage space of the host and storage system caches.
Such duplication is wasteful because of the
relatively high cost of implementing caches in the
host system and the storage systems.

004433-034950

SUMMARY OF THE INVENTION

5 A computer system is disclosed with mechanisms
for exclusive caching that avoid the accumulation of
duplicate copies of information in host and storage
system caches. A computer system according to these
exclusive caching techniques includes a host system
having a host cache and a storage system having a
storage system cache and functionality for performing
10 demote operations to coordinate the placement of
information in the host cache to the storage system
caches.

15 Other features and advantages of the present
invention will be apparent from the detailed
description that follows.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is described with respect to particular exemplary embodiments thereof and reference is accordingly made to the drawings in which:

Figure 1 shows a computer system that incorporates the exclusive caching techniques which are disclosed herein;

Figure 2 illustrates an eviction and demotion of a data block from a host cache according to the present techniques;

Figure 3 illustrates the handling of a demote operation by a storage system.

DETAILED DESCRIPTION

5 **Figure 1** shows a computer system 100 that
incorporates the exclusive caching techniques which
are disclosed herein. The computer system 100
includes a host system 10 and a storage system 20
that communicate via a communication path 30. The
host system 10 includes a host processor 12 and a
host cache 14. The storage system 20 includes a
10 storage system cache 22 and a storage media 24.

15 The host system 10 obtains information from the
storage system 20 by performing read operations via
the communication path 30. The host system 10
transfers information to the storage system 20 by
performing write operations via the communication
path 30. In addition, the host system 10 demotes
information from the host cache 14 to the storage
system cache 22 by performing demote operations via
20 the communication path 30. An example demote
operation 50 is shown which carries a demoted data
block 52 to the storage system 20 via the
communication path 30.

25 In one embodiment, the host cache 14 and the
storage system cache 22 are organized into sets of
information. These sets of information are
hereinafter referred to as data blocks but may also
be referred to as cache blocks, cache lines, or data
30 lines. A data block hereinafter refers to a block of
data having a particular address whether the data
block is stored in a cache or storage media.

The host processor 12 accesses data by performing a read or a write to the host cache 14. If the data referenced in the read or the write from the host processor 12 is not held in the host cache 14, then the host system 10 performs a read operation via the communication link 30 to obtain a data block that contains the data needed by the host processor 12 from the storage system 20.

10 In response to the read operation by the host system 10, the storage system 20 obtains the data block specified in the read operation from storage media 24 and transfers it to the host system 10 via the communication link 30. In embodiments of the
15 storage system 20 that perform read-ahead operations from the storage media 24, the data block specified in the read operation may be in the storage system cache 22. If so, the storage system 20 obtains the data block from storage system cache 22 and transfers
20 it to the host system 10 via the communication link 30 and then evicts the data block from the storage system cache 22.

The host system 10 receives the data block from the storage system 20 via the communication link 30 and stores the data block into the host cache 14. The host system 10 then completes the read or write operation for the host processor 12.

30 The process of storing a new data block into the host cache 14 may cause an eviction of another data block from the host cache 14. If the evicted data block is clean, then rather than being discarded as

in prior systems it is instead demoted to the storage system cache 22. An evicted data block is demoted by transferring it to the storage system 20 via the communication link 30 using a demote operation. For example, the demote operation 50 carries the demoted data block 52 which was a clean data block evicted from the host cache 14.

The storage system 20 handles the demote operation 50 by storing the demoted data block 52 into the storage system cache 22 while performing an eviction from the storage system cache 22 if needed. The initial placement of a data block into the host cache 14 and the eventual eviction and demotion of that data block to the storage system cache 22 avoids the creation of duplicate copies of the data block in the host cache 14 and the storage system cache 22.

The host system 10 includes a cache controller for managing the host cache 14. The cache controller in the host system 10 may implement any one or more of a set of known replacement policies for selecting data blocks to be evicted from the host cache 14. For example, the host system 10 may implement a least recently used (LRU) replacement policy for the host cache 14. Alternatively, the host system 10 may implement a most recently used (MRU) replacement policy for the host cache 14. In yet another alternative, the host system 10 may implement a most frequently replaced (MFR) replacement policy. These are only examples of suitable replacement policies and others may be employed.

004433-0310

The storage system 20 may be any type of storage system which is capable of holding data blocks for access by the host system 10 and which includes a storage system cache. In one embodiment, the storage media 24 is a disk drive. In another embodiment, the storage media 24 is an array of disk drives. In yet another embodiment, the storage media 24 is a solid-state memory. In another embodiment, the storage media 24 is a tape. The storage system 20 may implement the functionality of a server wherein the host system 10, as well as other hosts systems on the communication path 30, may access the storage system 10 using any one or more of a wide variety of known client-server communication protocols.

In an embodiment in which the storage media 24 is a tape, the host cache 14 may be a RAM cache or a disk cache.

The storage system cache 22 may be employed to hold data blocks which are obtained by read-ahead operations on the storage media 24 as well as to hold the data blocks which are demoted from the host system 10. In one embodiment, data blocks which are demoted from the host system 10 may be stored in any free area in the storage system cache 22. Alternatively, the storage of demoted data blocks may be limited to predetermined portions or sub-areas of the storage system cache 22. The storage system cache 22 may also be employed to hold other data blocks from read operations performed by the host system 10. The storage system cache 22 may be

employed to hold other data blocks from write operations performed by the host system 10.

5 The storage system 20 includes a cache controller that may implement any known data block replacement policy for the storage system cache 22.

10 In addition, the cache controller in the storage system 20 may preferentially keep data blocks in the storage system cache 22 that are being shared by multiple host systems that have access to the storage system 20. For example, the cache controller may select data blocks for eviction from among the unshared data blocks only and select shared data
15 blocks for eviction only if no unshared data blocks remain in the storage system cache 22. Alternatively, the cache controller may take into account the shared status of a data block as only one factor in its replacement policy with other factors
20 being the age, time of last update, etc., for the data block. These factors may be weighted in any combination.

25 The communication path 30 may be implemented as any type of communication path that is cable of carrying information between the host system 10 and the storage system 20. For example, the communication path 30 may be a system bus or a peripheral bus in a computer system. In another
30 example, the communication path 30 may be a network communication path. The communication path 30 may include a combination of bus and network elements. The host system 10 and the storage system 20 include

the appropriate communication hardware and software elements for performing communication via the particular embodiment of the communication path 30.

5 **Figure 2** illustrates an eviction and demotion of a data block from the host cache 14 according to the present techniques. At step 60, the cache controller in the host system 10 selects a data block and the host cache 14 for eviction. Any type of replacement
10 policy may be employed at step 60 to select the data block for eviction.

At step 62, the cache controller in the host system 10 determines whether the data block selected
15 at steps 60 is dirty. A dirty data block refers to a data block whose contents have changed since being entered into the host cache 14. If the data block selected at step 60 is dirty, then at step 64 the host system 10 performs a write operation to the
20 storage system 20 to write back the evicted data block to the storage system 20.

If the data block selected at step 60 is not dirty, then at step 66 the host system 10 performs a
25 demote operation for the evicted data block. The demote operation carries the evicted data block from the host cache 14 to the storage system 20 as a demoted data block. For example, the demote operation 50 carries the demoted data block 52.

30 **Figure 3** illustrates the handling of the demote operation 50 by the storage system 20. At step 70, the cache controller in the storage system 20

determines whether space is available in the storage system cache 22 to hold the demoted data block 52. If space is available in the storage system cache 22, then the demoted data block 52 is written into a free entry in the storage system cache 22 at step 72.

A free entry in the storage system cache 22 may not be available at step 70 if, for example, the storage system cache 22 has been filled by read-ahead operations from the storage media 24. In another example, a portion of the storage system cache 22 which is allocated to holding demoted data blocks may be filled with previously demoted data blocks.

If space is not available in the storage system cache 22 for the demoted data block 52 at step 70, then at step 74 an attempt is made to evict a data block from the storage system cache 22. If a data block was successfully evicted from the storage system cache 22 then at step the 80 the demoted data block 52 is written into the newly freed entry in the storage system cache 22. Otherwise, the demoted data block is discarded at step 78.

The foregoing detailed description of the present invention is provided for the purposes of illustration and is not intended to be exhaustive or to limit the invention to the precise embodiment disclosed. Accordingly, the scope of the present invention is defined by the appended claims.

CLAIMS

What is claimed is:

- 5 1. A method for exclusive caching in a computer
system, comprising the steps of:
 reading a set of information from a storage
system in the computer system;
 storing the information in a host cache in the
10 computer system;
 demoting the information from the host cache to
a storage system cache in the storage system.
- 15 2. The method of claim 1, wherein the step of
demoting comprises the steps of:
 evicting the information from the host cache;
 transferring the information to the storage
system in a demote operation on a communication path
between the host cache and the storage system.
- 20 3. The method of claim 2, wherein the step of
transferring the information to the storage system in
a demote operation is performed only if the
information is not dirty and the step of transferring
25 the information to the storage system in a write
operation is performed otherwise.
- 30 4. The method of claim 2, further comprising the
step of storing the information carried in the demote
operation into the storage system cache.
5. The method of claim 4, wherein the step of
storing the information carried in the demote

operation into the storage system cache comprises the steps of:

performing an eviction from the storage system cache;

5 discarding the information if the eviction fails.

6. The method of claim 5, wherein the step of performing an eviction from the storage system cache
10 comprises the step of selecting a data block for eviction from the storage system cache from among a set of unshared data blocks in the storage system cache.

15 7. The method of claim 5, wherein the step of performing an eviction from the storage system cache comprises the step of selecting a data block for eviction from the storage system cache using a shared status of the data blocks as a factor in a
20 replacement policy.

8. A computer system, comprising:
 storage system having a storage media and a storage system cache;
25 host system having a host cache;
 means for exclusively caching a set of information obtained from the storage media such that the information is stored either in the host cache or the storage system cache but not both caches at the
30 same time.

9. The computer system of claim 8, wherein the means for exclusively caching includes means for

0044333-00100

evicting the information from the host cache and transferring the information to the storage system in a demote operation on a communication path between the host cache and the storage system.

5

10. The computer system of claim 9, wherein the communication path is a bus in the computer system.

11. The computer system of claim 9, wherein the communication path is a network communication link in the computer system.

12. The computer system of claim 9, wherein the means for exclusively caching further includes means for storing the information carried in the demote operation into the storage system cache.

13. The computer system of claim 12, wherein the means for storing the information carried in the demote operation into the storage system cache includes means for performing an eviction from the storage system cache and discarding the information if the eviction fails.

14. The computer system of claim 13, wherein the means for performing an eviction from the storage system cache comprises means for selecting a data block for eviction from the storage system cache from among a set of unshared data blocks in the storage system cache.

15. The computer system of claim 13, wherein the means for performing an eviction from the storage

system cache comprises means for selecting a data block for eviction from the storage system cache using a shared status of the data blocks as a factor in a replacement policy.

5

16. The computer system of claim 8, wherein the storage system cache includes a sub-area for exclusive caching and sub-area holding other cached information.

10

17. The computer system of claim 16, wherein the other cached information is obtained by performing read-ahead operations from the storage media.

15

18. The computer system of claim 8, wherein the storage media is a disk array.

19. The computer system of claim 8, wherein the storage media is a tape.

A computer system with mechanisms for exclusive caching that avoids the accumulation of duplicate copies of information in host and storage system caches. A computer system according to these exclusive caching techniques includes a host system having a host cache and a storage system having a storage system cache and functionality for performing demote operations to coordinate the placement of information in the host cache to the storage system caches.




FIGURE 1

10004565

```

graph TD
    Start([Demote Operation]) --> Decision70{Storage System Cache Avail.  
?}
    Decision70 -- Y --> Process72[Write the Demoted Data Block Into a  
Free Entry in the Storage System Cache]
    Process72 --> Done1([Done])
    Decision70 -- N --> Process74[Attempt to Evict a Data Block  
From the Storage System Cache]
    Process74 --> Decision76{Success ?}
    Decision76 -- Y --> Process80[Write the Demoted Data Block Into a  
Free Entry in the Storage System Cache]
    Process80 --> Done2([Done])
    Decision76 -- N --> Process78[Discard the Demoted Data Block]
    Process78 --> Done3([Done])
  
```

10004565

**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION**ATTORNEY DOCKET NO. 10004565

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

EXCLUSIVE CACHING IN COMPUTER SYSTEMS

the specification of which is attached hereto unless the following box is checked:

() was filed on _____ as US Application Serial No. or PCT International Application
Number _____ and was amended on _____ (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
			YES: _____ NO: _____
			YES: _____ NO: _____

Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION SERIAL NUMBER	FILING DATE

U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, I acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION SERIAL NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number **022879**Place Customer
Number Bar Code
Label here

Send Correspondence to:
HEWLETT-PACKARD COMPANY
 Intellectual Property Administration
 P.O. Box 272400
 Fort Collins, Colorado 80527-2400

Direct Telephone Calls To:

Paul H. Horstmann
(415) 602-1721

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: John WilkesCitizenship: United KingdomResidence: 3788 Park Blvd. Palo Alto, CA 94306Post Office Address: 3788 Park Blvd. Palo Alto, CA 94306Inventor's Signature John WilkesDate 8 August 2000


**DECLARATION AND POWER OF ATTORNEY
FOR PATENT APPLICATION (continued)**

ATTORNEY DOCKET NO. 10004565

Full Name of # 2 joint inventor: Theodore M. Wong Citizenship: U.S.

Residence: 1635 Denniston Avenue Pittsburgh, PA 15217

Post Office Address: 1635 Denniston Avenue Pittsburgh, PA 15217

Inventor's Signature  Date 8th Aug 2000

Full Name of # 3 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 4 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 5 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 6 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 7 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____

Full Name of # 8 joint inventor: _____ Citizenship: _____

Residence: _____

Post Office Address: _____

Inventor's Signature _____ Date _____